Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	1509	711/103.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:17
L2	1010	711/105.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:17
L3	27187	"711"/$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:17
L4	832	(Yoshiyuki near Tanaka)".in"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:17
L5	7	(Makoto near Yatabe).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:18
L6	34	(Takeaki near Sato).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:18
L7	11	(Kazuya near Kawamoto).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:19
L8	861	4 or 5 or 6 or 7	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:20
L9	5185	(translation adj lookaside adj buffer) or TLB	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:21
L10	6962	(translation near2 table)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:21

L11	11399	9 or 10	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:21
L12	240927	memory near3 (flash or nonvolatile or (non adj volatile) or programmable or $2prom)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:22
L13	96	first adj memory adj zone	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:23
L14	77	second adj memory adj zone	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:23
L15	21771	block near3 pair	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:23
L16	1828	11 and 12	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:23
L17	34	16 and 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:23
L18	60	13 and 14	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:23
L19	0	17 and 18	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:24
L20	4	3 and 18	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:24

L21	2	20 and 12	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 16:24

Terms used
translation lookaside buffer TLB flash flash programmable memory eeprom

Found 17 of 167,655

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 17 of 17

Relevance scale

1 SoftFLASH: analyzing the performance of clustered distributed virtual shared memory

Andrew Erlichson, Neal Nuckolls, Greg Chesson, John Hennessy

September 1996 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the seventh international conference on Architectural support for programming languages and operating systems ASPLOS-VII**, Volume 31 , 30 Issue 9 , 5

Publisher: ACM Press

Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, references, citings, index terms

One potentially attractive way to build large-scale shared-memory machines is to use small-scale to medium-scale shared-memory machines as clusters that are interconnected with an off-the-shelf network. To create a shared-memory programming environment across the clusters, it is possible to use a virtual shared-memory software layer. Because of the low latency and high bandwidth of the interconnect available within each cluster, there are clear advantages in making the clusters as large as possi ...

2 Disco: running commodity operating systems on scalable multiprocessors

Edouard Bugnion, Scott Devine, Kinshuk Govil, Mendel Rosenblum

November 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 4

Publisher: ACM Press

Full text available: pdf(400.76 KB) Additional Information: full citation, abstract, references, citings, index terms, review

In this article we examine the problem of extending modern operating systems to run efficiently on large-scale shared-memory multiprocessors without a large implementation effort. Our approach brings back an idea popular in the 1970s: virtual machine monitors. We use virtual machines to run multiple commodity operating systems on a scalable multiprocessor. This solution addresses many of the challenges facing the system software for these machines. We demonstrate our approach with a prototy ...

Keywords: scalable multiprocessors, virtual machines

3 Disco: running commodity operating systems on scalable multiprocessors

Edouard Bugnion, Scott Devine, Mendel Rosenblum

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: pdf(2.30 MB) Additional Information: full citation, references, citings, index terms

4 Decoupled hardware support for distributed shared memory

Steven K. Reinhardt, Robert W. Pfile, David A. Wood

May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture ISCA '96**, Volume 24 Issue 2

Publisher: ACM Press

Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper investigates hardware support for fine-grain distributed shared memory (DSM) in networks of workstations. To reduce design time and implementation cost relative to dedicated DSM systems, we decouple the functional hardware components of DSM support, allowing greater use of off-the-shelf devices.We present two decoupled systems, Typhoon-0 and Typhoon-1. Typhoon-0 uses an off-the-shelf protocol processor and network interface; a custom access control device is the only DSM-specific hard ...

5 Options for dynamic address translation in COMAs

Xiaogang Qiu, Michel Dubois

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.37 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

In modern processors, the dynamic translation of virtual addresses to support virtual memory is done before or in parallel with the first-level cache access. As processor technology improves at a rapid pace and the working sets of new applications grow insatiably the latency and bandwidth demands on the TLB (Translation Lookaside Buffer) are getting more and more difficult to meet. The situation is worse in multiprocessor systems, which run larger applications and are plagued by the TLB consiste ...

6 Multigrain shared memory

Donald Yeung, John Kubiatowicz, Anant Agarwal

May 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 2

Publisher: ACM Press

Full text available: pdf(369.18 KB) Additional Information: full citation, abstract, references, index terms, review

Parallel workstations, each comprising tens of processors based on shared memory, promise cost-effective scalable multiprocessing. This article explores the coupling of such small- to medium-scale shared-memory multiprocessors through software over a local area network to synthesize larger shared-memory systems. We call these systems Distributed Shared-memory MultiProcessors (DSMPs). This article introduces the design of a shared-memory system that uses multiple granularities of sharing, ca ...

Keywords: distributed memory, symmetric multiprocessors, system of systems

7

Synchronization and communication in the T3E multiprocessor
Steven L. Scott
September 1996 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the seventh international conference on Architectural support for programming languages and operating systems ASPLOS-VII**, Volume 31 , 30 Issue 9 , 5
Publisher: ACM Press
Full text available: pdf(1.34 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the synchronization and communication primitives of the Cray T3E multiprocessor, a shared memory system scalable to 2048 processors. We discuss what we have learned from the T3D project (the predecessor to the T3E) and the rationale behind changes made for the T3E. We include performance measurements for various aspects of communication and synchronization.The T3E augments the memory interface of the DEC 21164 microprocessor with a large set of explicitly-managed, external r ...

8 Cashmere-2L: software coherent shared memory on a clustered remote-write network
Robert Stets, Sandhya Dwarkadas, Nikolaos Hardavellas, Galen Hunt, Leonidas Kontothanassis, Srinivasan Parthasarathy, Michael Scott
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5
Publisher: ACM Press
Full text available: pdf(2.17 MB) Additional Information: full citation, references, citings, index terms

9 Tempest and typhoon: user-level shared memory
S. K. Reinhardt, J. R. Larus, D. A. Wood
April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST annual international symposium on Computer architecture ISCA '94**, Volume 22 Issue 2
Publisher: IEEE Computer Society Press, ACM Press
Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, references, citings, index terms

Future parallel computers must efficiently execute not only hand-coded applications but also programs written in high-level, parallel programming languages. Today's machines limit these programs to a single communication paradigm, either message-passing or shared-memory, which results in uneven performance. This paper addresses this problem by defining an interface, *Tempest*, that exposes low-level communication and memory-system mechanisms so programmers and compilers can customize polici ...

10 The M-Machine multicomputer
Marco Fillo, Stephen W. Keckler, William J. Dally, Nicholas P. Carter, Andrew Chang, Yevgeny Gurevich, Whay S. Lee
December 1995 **Proceedings of the 28th annual international symposium on Microarchitecture**
Publisher: IEEE Computer Society Press
Full text available: pdf(1.29 MB) Additional Information: full citation, references, citings, index terms

11 Tempest and typhoon: user-level shared memory
Steven K. Reinhardt, James R. Larus, David A. Wood
August 1998

25 years of the international symposia on Computer architecture (selected papers)
Publisher: ACM Press
Full text available: pdf(1.57 MB) Additional Information: full citation, references, index terms

12 AP1000+: architectural support of PUT/GET interface for parallelizing compiler
Kenichi Hayashi, Tsunehisa Doi, Takeshi Horie, Yoichi Koyanagi, Osamu Shiraki, Nobutaka Imamura, Toshiyuki Shimizu, Hiroaki Ishihata, Tatsuya Shindo
November 1994 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the sixth international conference on Architectural support for programming languages and operating systems ASPLOS-VI**, Volume 29 , 28 Issue 11 , 5
Publisher: ACM Press
Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references, citings, index terms

The scalability of distributed-memory parallel computers makes them attractive candidates for solving large-scale problems. New languages, such as HPF, FortranD, and VPP Fortran, have been developed to enable existing software to be easily ported to such machines. Many distributed-memory parallel computers have been built, but none of them support the mechanisms required by such languages. We studied the mechanisms required by parallelizing compilers and proposed a new architecture to suppo ...

13 Operating system support for high-speed communication
Peter Druschel
September 1996 **Communications of the ACM**, Volume 39 Issue 9
Publisher: ACM Press
Full text available: pdf(313.01 KB) Additional Information: full citation, references, citings, index terms, review

14 Using the SimOS machine simulator to study complex computer systems
Mendel Rosenblum, Edouard Bugnion, Scott Devine, Stephen A. Herrod
January 1997 **ACM Transactions on Modeling and Computer Simulation (TOMACS)**, Volume 7 Issue 1
Publisher: ACM Press
Full text available: pdf(731.76 KB) Additional Information: full citation, references, citings, index terms, review

Keywords: computer architecture, computer simulation, computer system performance analysis, operating systems

15 Compiler-directed page coloring for multiprocessors
Edouard Bugnion, Jennifer M. Anderson, Todd C. Mowry, Mendel Rosenblum, Monica S. Lam
September 1996 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the seventh international conference on Architectural support for programming languages and operating systems ASPLOS-VII**, Volume 31 , 30 Issue 9 , 5
Publisher: ACM Press
Full text available: pdf(1.37 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents a new technique, *compiler-directed page coloring,* that eliminates conflict misses in multiprocessor applications. It enables applications to make better use of the increased aggregate cache size available in a multiprocessor. This technique uses the compiler's knowledge of the access patterns of the parallelized applications to direct the operating system's virtual memory page mapping strategy. We demonstrate that this technique can lead to significant performance impr ...

16 Performance analysis using the MIPS R10000 performance counters

Marco Zagha, Brond Larson, Steve Turner, Marty Itzkowitz

November 1996 **Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM) - Volume 00 Supercomputing '96**

Publisher: ACM Press, IEEE Computer Society

Full text available: pdf(200.39 KB) Additional Information: full citation, abstract, references, citings, index terms

Tuning supercomputer application performance often requires analyzing the interaction of the application and the underlying architecture. In this paper, we describe support in the MIPS R10000 for non-intrusively monitoring a variety of processor events -- support that is particularly useful for characterizing the dynamic behavior of multi-level memory hierarchies, hardware-based cache coherence, and speculative execution. We first explain how performance data is collected using an integrate ...

Keywords: performance analysis, profiling tools, hardware performance counters, MIPS R10000, SGI Power Challenge

17 Evaluating the impact of simultaneous multithreading on network servers using real hardware

Yaoping Ruan, Vivek S. Pai, Erich Nahum, John M. Tracey

June 2005 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '05**, Volume 33 Issue 1

Publisher: ACM Press

Full text available: pdf(498.70 KB) Additional Information: full citation, abstract, references, index terms

This paper examines the performance of simultaneous multithreading (SMT) for network servers using actual hardware, multiple network server applications, and several workloads. Using three versions of the Intel Xeon processor with Hyper-Threading, we perform macroscopic analysis as well as microarchitectural measurements to understand the origins of the performance bottlenecks for SMT processors in these environments. The results of our evaluation suggest that the current SMT support in the Xeon ...

Keywords: network server, simultaneous multithreading(SMT)

Results 1 - 17 of 17

Wed, 7 Dec 2005, 4:31:46 PM EST

Edit an existing query or compose a new query in the Search Query Display.

Search Query Display

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

#1 ((translation lookaside buffer)<in>metadata)

#2 ((translation lookaside buffer)<in>metadata) or tlb

#3 (memory zone<IN>metadata)

#4 (block <IN>metadata)

#5 (pair<IN>metadata)

#6 (programmable memory<IN>metadata)

#7 (flash or eeprom or eprom or rom or prom<IN>metadata)

#8 (translation table<IN>metadata)

#9 (translation buffer<IN>metadata)